arXiv:1605.01942v1 [math.CO] 6 May 2016

EDGE-COVERS IN d-INTERVAL HYPERGRAPHS

RON AHARONI, RON HOLZMAN, AND SHIRA ZERBIB

Abstract. A d-interval hypergraph has d disjoint copies of the unit interval as its vertex set, and each edge is the union of d subintervals, one on each copy. Extending a classical result of Gallai on the case $d = 1$, Tardos and Kaiser used topological tools to bound the ratio between the transversal number and the matching number in such hypergraphs. We take a dual point of view, and bound the edge-covering number (namely the minimal number of edges covering the entire vertex set) in terms of a parameter expressing independence of systems of partitions of the d unit intervals. The main tool we use is an extension of the KKM theorem to products of simplices, due to Peleg. Our approach also yields a new proof of the Tardos-Kaiser result.

1. Introduction

A *d-interval hypergraph* has vertex set $V = \bigcup_{i=1}^{d} U^i$, where the U^i are d disjoint copies of the unit interval $[0, 1]$. It has a finite edge set E, each edge being of the form $e = \bigcup_{i=1}^{d} I^i$ where I^i is a non-empty closed subinterval of U^i. In the case $d = 1$ we get the familiar interval hypergraphs. There are other variants of d-interval hypergraphs in the literature [7, 8, 2, 3, 1], including those where edges consist of unions of d intervals on the same copy of $[0, 1]$, and discrete versions where the vertex set is a finite ordered set; these will not be discussed here.

We recall that in a hypergraph $H = (V, E)$, a *matching* is a set of disjoint edges, and $\nu(H)$ denotes the *matching number* – the maximal size of a matching. A *transversal* (also called *vertex-cover*) is a set of vertices meeting all edges, and $\tau(H)$ denotes the *transversal number* – the minimal size of a transversal. We always have $\nu(H) \leq \tau(H)$.

By reversing the roles of vertices and edges (duality of hypergraphs), we get two corresponding concepts. Namely, in a hypergraph $H = (V, E)$, a set of vertices is *strongly independent* if no two of them belong to an edge. (An independent set is one that does not contain any edge, hence the "strongly" in our terminology.) We write $\iota(H)$ for the *strong independence number* – the maximal size of a strongly independent set ($\iota(H) = \infty$ is allowed). A

The research of the first author was supported by BSF grant no. 2006099, by an ISF grant and by the Discount Bank Chair at the Technion.

The research of the third author was partly supported by the New-England Fund, Technion.

set of edges is an *edge-cover* if their union is V. We write $\rho(H)$ for the *edge-covering number* – the minimal size of an edge-cover ($\rho(H) = \infty$ if $\cup E \neq V$). We always have $\iota(H) \leq \rho(H)$.

It is well known that for interval hypergraphs equality holds for each of these pairs of parameters (part (a) below is due to Gallai [6]):

Theorem 1.1. *Let $H = (V, E)$ be an interval hypergraph. Then:*

(a) $\nu(H) = \tau(H)$.
(b) $\iota(H) = \rho(H)$.

Proof. We briefly recall the standard constructive arguments.

(a) Let $I = [a, b]$ be an interval in E with the leftmost right endpoint. Then we can place I in the matching being constructed, b in the transversal, and proceed by induction with the subhypergraph consisting of those intervals in E having their left endpoint to the right of b.
(b) We may assume that $\cup E = V$, otherwise $\iota(H) = \rho(H) = \infty$. Let $I = [0, c]$ be an interval in E with the rightmost right endpoint among those starting at 0. Then we can place 0 in the strongly independent set being constructed, I in the edge-cover, and proceed by induction with the hypergraph having as ground set $[c + \varepsilon, 1]$ for small enough $\varepsilon > 0$, and as edges the non-empty intersections of intervals in E with the ground set.

□

For d-interval hypergraphs with $d \geq 2$ these equalities need not hold, but one can bound the ratio between the corresponding pairs of parameters. In the case of ν and τ this was a challenging problem, solved by Tardos [15] for $d = 2$ and Kaiser [8] for general d:

Theorem 1.2 ([15, 8])**.** *Let H be a d-interval hypergraph, $d \geq 2$. Then $\tau(H) \leq d(d-1)\nu(H)$.*

Both Tardos and Kaiser used topological methods. Alon [2] obtained a slightly weaker bound by combinatorial means.

In the case of ι and ρ, bounding the ratio for d-interval hypergraphs is easy:

Observation 1.3. *Let H be a d-interval hypergraph. Then $\rho(H) \leq d\iota(H)$, and this bound is tight.*

Proof. For $i = 1, \ldots, d$, let H^i be the interval hypergraph on U^i whose edges are the intersections of the edges of H with U^i. Clearly $\iota(H^i) \leq \iota(H)$, and so by Theorem 1.1(b) we have $\rho(H^i) \leq \iota(H)$, $i = 1, \ldots, d$. Together, these d edge-covers yield $\rho(H) \leq d\iota(H)$.

To see that this is tight, fix d and n. We construct a d-interval hypergraph $H = (V, E)$ with $\iota(H) = n$ and $\rho(H) = dn$. For each $i = 1, \ldots, d$, we use two systems of subintervals of U^i. The first consists of $J^i_j = [\frac{j-1}{n}, \frac{j}{n}]$,

$j = 1, \ldots, n$; the second consists of $K^i_k = [\frac{k-1}{(dn)^2}, \frac{k}{(dn)^2}]$, $k = 1, \ldots, (dn)^2$. An edge $e \in E$ is determined by a choice of $i \in \{1, \ldots, d\}$, $j \in \{1, \ldots, n\}$, and $k_{i'} \in \{1, \ldots, (dn)^2\}$ for each $i' \in \{1, \ldots, d\} \setminus \{i\}$. The corresponding edge is $e = J^i_j \cup \bigcup_{i' \neq i} K^{i'}_{k_{i'}}$. Clearly, no strongly independent set can contain points in two distinct $U^i, U^{i'}$, and it can contain at most n points in the same U^i; hence $\iota(H) = n$. The total length of an edge $e \in E$ is $\frac{1}{n} + \frac{d-1}{(dn)^2} = \frac{d^2n+d-1}{(dn)^2}$, and hence the total length of $dn-1$ such edges is $\frac{(dn-1)(d^2n+d-1)}{(dn)^2} < d$. Thus, there is no edge-cover of size $dn-1$, and hence $\rho(H) = dn$. □

There is, however, another way of looking at Theorem 1.1, which will lead to a more interesting extension of part (b) to the case of d-interval hypergraphs.

A system of n subintervals of $[0, 1]$ of the form

$$[0, c_1), (c_1, c_2), \ldots, (c_{n-2}, c_{n-1}), (c_{n-1}, 1],$$

where $0 \leq c_1 \leq c_2 \leq \cdots \leq c_{n-2} \leq c_{n-1} \leq 1$, will be called an *$n$-partition* of $[0, 1]$. The subintervals will be referred to as the first, the second, ..., the $(n-1)$-th, the n-th *cell* of the partition. Note that the j-th cell is empty if $c_{j-1} = c_j$ (with $c_0 = 0, c_n = 1$).

Theorem 1.1 can be equivalently re-stated as:

Theorem 1.4. *Let $H = (V, E)$ be an interval hypergraph, and let n be a positive integer. Then:*

(a) *If every n-partition of $[0, 1]$ has a cell that contains an edge of H, then there exists an n-partition all of whose cells contain edges of H.*
(b) *If every n-partition of $[0, 1]$ has a cell that is contained in an edge of H, then there exists an n-partition all of whose cells are contained in edges of H.*

Now we generalize the partition concept to the setting where $V = \bigcup_{i=1}^d U^i$. A *$d \times n$-partition* of V consists of d n-partitions of $U^1, \ldots, U^d$ respectively. Its total number of cells is dn. Any union of d cells, one from each U^i, will be called a *d-cell*. Thus, there are n^d d-cells.

Part (a) of Theorem 1.4 has the following d-counterpart:

Theorem 1.5. *Let $H = (V, E)$ be a d-interval hypergraph, $d \geq 2$, and let n be a positive integer. If every $d \times n$-partition of V has a d-cell that contains an edge of H, then there exists a $d \times n$-partition having at least $\frac{n}{d-1}$ disjoint d-cells that contain edges of H.*

Essentially, Theorem 1.2 was proved in [15, 8] by establishing Theorem 1.5, and then applying it with the largest n such that $d(n-1) < \tau(H)$.

Our main result is the following d-counterpart of Theorem 1.4(b):

Theorem 1.6. *Let $H = (V, E)$ be a d-interval hypergraph, and let n be a positive integer. Assume that every $d \times n$-partition of V has a d-cell that is contained in an edge of H. Then there exists a $d \times n$-partition all of whose*

cells are contained in edges of H_0, where H_0 is a subhypergraph of H of size at most $(1+\ln d)n$. In particular, $\rho(H) \leq (1+\ln d)n$.
Moreover, for $d = 2$ the upper bound on the size of H_0 can be improved to n, and hence $\rho(H) \leq n$ for $d = 2$.

Although Theorem 1.5 is not new, we give here a new proof of it, which goes along similar lines to our proof of Theorem 1.6. Both proofs are topological, using extensions due to Peleg [12, 13] of the well-known KKM theorem [9].

In Section 2 we recall the KKM theorem about coverings of a simplex, and a dual variant due to Sperner [14]. Then, as a warm-up to our main proofs, we apply them to get topological proofs of both parts of Theorem 1.4 (and thus of Theorem 1.1). Then, in Section 3, we recall the extensions due to Peleg of the KKM/Sperner theorems to coverings of a product of simplices. We also recall some well-known bounds on the ratio between integer and fractional versions of the matching and transversal numbers of hypergraphs. Equipped with these tools, we prove Theorems 1.5 and 1.6 in Section 4. In the final Section 5, we discuss the tightness of the bound in Theorem 1.6. We show that $\rho(H) \leq n$ is best possible for $d = 2$, and that it does not hold anymore for $d > 2$. The $(1+\ln d)n$ upper bound for general d can be somewhat improved, but we do not know by how much.

2. Topological proofs for interval hypergraphs

We denote by Δ_{n-1} the standard simplex in $\mathbb{R}^n$, namely

$$\Delta_{n-1} = \{\vec{x} = (x_1, \ldots, x_n) \mid \sum_{j=1}^{n} x_j = 1, x_j \geq 0, j = 1, \ldots, n\}.$$

For $S \subseteq [n] = \{1, \ldots, n\}$, $S \neq \emptyset$, we denote by $F(S)$ the face of Δ_{n-1} spanned by the corresponding unit vectors, i.e.,

$$F(S) = \{\vec{x} = (x_1, \ldots, x_n) \in \Delta_{n-1} \mid x_j = 0 \text{ for all } j \notin S\}.$$

Theorem 2.1 (KKM [9])**.** *Let $A_1, \ldots, A_n$ be subsets of Δ_{n-1} that are all closed or all open. Suppose that for every $\emptyset \neq S \subseteq [n]$ we have $F(S) \subseteq \bigcup_{j \in S} A_j$. Then $\bigcap_{j=1}^{n} A_j \neq \emptyset$.*

We shall use the following immediate corollary of the KKM theorem:

Corollary 2.2. *Let $A_1, \ldots, A_n$ be subsets of Δ_{n-1} that are all closed or all open. Suppose that*

(a) *$\bigcup_{j=1}^{n} A_j = \Delta_{n-1}$, and*
(b) *for all $\vec{x} = (x_1, \ldots, x_n) \in \Delta_{n-1}$ and all $j \in [n]$, $x_j = 0 \Rightarrow \vec{x} \notin A_j$.*

Then $\bigcap_{j=1}^{n} A_j \neq \emptyset$.

There is a dual variant which differs from Corollary 2.2 by reversing the conclusion of $x_j = 0$:

Theorem 2.3 (Sperner [14])**.** *Let $B_1, \ldots, B_n$ be subsets of Δ_{n-1} that are all closed or all open. Suppose that*

(a) $\bigcup_{j=1}^{n} B_j = \Delta_{n-1}$*, and*

(b) *for all $\vec{x} = (x_1, \ldots, x_n) \in \Delta_{n-1}$ and all $j \in [n]$, $x_j = 0 \Rightarrow \vec{x} \in B_j$.*

Then $\bigcap_{j=1}^{n} B_j \neq \emptyset$.

It is possible to derive Theorem 2.3 from Corollary 2.2 by letting $A_j = \Delta_{n-1} \setminus B_j$.

There is a canonical bijection between n-partitions of $[0,1]$, as defined in the Introduction, and points in Δ_{n-1}. Namely, the n-partition

$$[0, c_1), (c_1, c_2), \ldots, (c_{n-2}, c_{n-1}), (c_{n-1}, 1]$$

corresponds to the point $\vec{x} = (x_1, x_2, \ldots, x_{n-1}, x_n)$, where $x_j = c_j - c_{j-1}$ is the length of the j-th cell (with $c_0 = 0, c_n = 1$). In the following, we identify every n-partition of $[0,1]$ with the corresponding $\vec{x} \in \Delta_{n-1}$, and use the topological results above to prove Theorem 1.4.

Proof of Theorem 1.4(a) from Corollary 2.2: Given the interval hypergraph $H = (V, E)$ and the positive integer n, we define $A_j \subseteq \Delta_{n-1}$ for $j \in [n]$ by:

$$\vec{x} \in A_j \;\Leftrightarrow\; \text{the } j\text{-th cell of the } n\text{-partition } \vec{x} \text{ contains an edge of } H$$

It is easy to check that the sets $A_1, \ldots, A_n$ are open, and that the premise of Theorem 1.4(a) renders them a covering of Δ_{n-1}. Moreover, condition (b) of Corollary 2.2 holds, since $x_j = 0$ means that the j-th cell is empty. It follows that there exists $\vec{x} \in \bigcap_{j=1}^{n} A_j$. Then, all cells of the n-partition corresponding to $\vec{x}$ contain edges of H, as desired. □

Proof of Theorem 1.4(b) from Theorem 2.3: Given the interval hypergraph $H = (V, E)$ and the positive integer n, we define $B_j \subseteq \Delta_{n-1}$ for $j \in [n]$ by:

$$\vec{x} \in B_j \;\Leftrightarrow\; \text{the } j\text{-th cell of the } n\text{-partition } \vec{x} \text{ is contained in an edge of } H$$

It is easy to check that the sets $B_1, \ldots, B_n$ are closed, and that the premise of Theorem 1.4(b) renders them a covering of Δ_{n-1}. Moreover, condition (b) of Theorem 2.3 holds, since an empty cell is contained in any edge. It follows that there exists $\vec{x} \in \bigcap_{j=1}^{n} B_j$, meaning that all cells of the corresponding n-partition are contained in edges of H. □

3. Tools for handling d-interval hypergraphs

We denote by $P_{d,n}$ the Cartesian product of d copies of Δ_{n-1}, i.e.,

$$P_{d,n} = \Delta_{n-1} \times \Delta_{n-1} \times \cdots \times \Delta_{n-1} = (\Delta_{n-1})^d.$$

For $i \in [d]$ and $\emptyset \neq S \subseteq [n]$ we denote by $P_{d,n}(i, S)$ the polytope obtained from $P_{d,n}$ by replacing the i-th factor in the product by its face $F(S)$, namely

$$P_{d,n}(i, S) = \Delta_{n-1} \times \cdots \times F(S) \times \cdots \times \Delta_{n-1}.$$

Theorem 3.1 (Peleg [12])**.** *Let A^i_j, $i = 1, \ldots, d$, $j = 1, \ldots, n$, be closed subsets of $P_{d,n}$. Suppose that for every $i \in [d]$ and $\emptyset \neq S \subseteq [n]$ we have $P_{d,n}(i, S) \subseteq \bigcup_{j \in S} A^i_j$. Then $\bigcap_{i=1}^d \bigcap_{j=1}^n A^i_j \neq \emptyset$.*

This is a generalization of the KKM theorem to d-fold products of simplices, and it likewise has the following immediate corollary:

Corollary 3.2. *Let A^i_j, $i = 1, \ldots, d$, $j = 1, \ldots, n$, be closed subsets of $P_{d,n}$. Suppose that*

(a) *for all $i \in [d]$, $\bigcup_{j=1}^n A^i_j = P_{d,n}$, and*
(b) *for all $\vec{x} = (x^1_1, \ldots, x^1_n; \cdots; x^d_1, \ldots, x^d_n) \in P_{d,n}$, all $i \in [d]$ and all $j \in [n]$, $x^i_j = 0 \Rightarrow \vec{x} \notin A^i_j$.*

Then $\bigcap_{i=1}^d \bigcap_{j=1}^n A^i_j \neq \emptyset$.

Here, too, there is a dual variant which differs from Corollary 3.2 by reversing the conclusion of $x^i_j = 0$:

Theorem 3.3 (Peleg [13])**.** *Let B^i_j, $i = 1, \ldots, d$, $j = 1, \ldots, n$, be closed subsets of $P_{d,n}$. Suppose that*

(a) *for all $i \in [d]$, $\bigcup_{j=1}^n B^i_j = P_{d,n}$, and*
(b) *for all $\vec{x} = (x^1_1, \ldots, x^1_n; \cdots; x^d_1, \ldots, x^d_n) \in P_{d,n}$, all $i \in [d]$ and all $j \in [n]$, $x^i_j = 0 \Rightarrow \vec{x} \in B^i_j$.*

Then $\bigcap_{i=1}^d \bigcap_{j=1}^n B^i_j \neq \emptyset$.

We remark that, unlike the case $d = 1$, there seems to be no direct way to deduce Theorem 3.3 from Corollary 3.2. Another comment is that Peleg's theorems hold also for products of simplices of different dimensions, but were stated above in the special case (needed here) where the simplices have the same dimension.

The other known results that we need concern hypergraphs $H = (V, E)$ on finite vertex sets, and involve fractional versions of our hypergraph invariants. We recall that a *fractional matching* is a function $f : E \to \mathbb{R}_+$ such that $\sum_{e \ni v} f(e) \leq 1$ for all $v \in V$. The *fractional matching number* $\nu^*(H)$ is the maximum of $\sum_{e \in E} f(e)$ over all fractional matchings in H. Similarly, a *fractional transversal* is a function $g : V \to \mathbb{R}_+$ such that $\sum_{v \in e} g(v) \geq 1$ for all $e \in E$. The *fractional transversal number* $\tau^*(H)$ is the minimum of $\sum_{v \in V} g(v)$ over all fractional transversals in H. By linear programming duality we always have $\nu^*(H) = \tau^*(H)$.

Analogously, a *fractional strongly independent set* is a function $f : V \to \mathbb{R}_+$ such that $\sum_{v \in e} f(v) \leq 1$ for all $e \in E$. The *fractional strong independence number* $\iota^*(H)$ is the maximum of $\sum_{v \in V} f(v)$ over all fractional strongly independent sets in H. A *fractional edge-cover* is a function $g : E \to \mathbb{R}_+$ such that $\sum_{e \ni v} g(e) \geq 1$ for all $v \in V$. The *fractional edge-covering number* $\rho^*(H)$ is the minimum of $\sum_{e \in E} g(e)$ over all fractional edge-covers in H. We always have $\iota^*(H) = \rho^*(H)$; both are defined as ∞ if $\cup E \neq V$.

We further recall that a hypergraph $H = (V, E)$ is *d-partite* if there exists a partition $V^1, \ldots, V^d$ of V such that $|e \cap V^i| = 1$ for every $e \in E$ and $i \in [d]$. We will use the following bound on the ratio ν^*/ν:

Theorem 3.4 (Furedi [4])**.** *Let H be a d-partite hypergraph, $d \geq 2$. Then $\nu(H) \geq \frac{\nu^*(H)}{d-1}$.*

We recall that the degree of a vertex v is the number of edges e that contain v. We denote by $\Delta(H)$ the *maximal degree* in H. Similarly, the size of an edge e is the number of vertices v contained in e. We denote by $r(H)$ the *rank* of H, i.e., the maximal size of an edge. A well-known bound on the ratio τ/τ^* is:

Theorem 3.5 (Lovász [10])**.** *Let H be a hypergraph with $\Delta(H) = d \geq 1$. Then $\tau(H) \leq (1 + \ln d)\tau^*(H)$.*

We will use this bound with the roles of vertices and edges reversed, namely:

Corollary 3.6. *Let H be a hypergraph with $r(H) = d \geq 1$. Then $\rho(H) \leq (1 + \ln d)\rho^*(H)$.*

In the case $d = 2$, when H is bipartite, the situation is even better:

Theorem 3.7 (Gallai [5])**.** *Let G be a bipartite graph. Then $\rho(G) = \rho^*(G)$.*

4. Proof of Theorems 1.5 and 1.6

Extending what we did in Section 2 in a natural way, we now identify every $d \times n$-partition of $V = \bigcup_{i=1}^d U^i$ with a point $\vec{x} = (x_1^1, \ldots, x_n^1; \cdots; x_1^d, \ldots, x_n^d)$ in $P_{d,n} = (\Delta_{n-1})^d$. Under this identification, x_j^i is the length of the j-th cell in the n-partition of U^i.

We use the set $[n]^d$ to index the d-cells of a $d \times n$-partition of V. That is, a d-tuple $\vec{j} = (j_1, \ldots, j_d) \in [n]^d$ will index the d-cell formed by the j_1-th cell in the n-partition of U^1, ..., the j_d-th cell in the n-partition of U^d. We refer to this as the $\vec{j}$-th d-cell.

Proof of Theorem 1.5: Given the d-interval hypergraph $H = (V, E)$, $d \geq 2$, and the positive integer n, we define $Y_{\vec{j}} \subseteq P_{d,n}$ for $\vec{j} \in [n]^d$ by:

$$\vec{x} \in Y_{\vec{j}} \;\Leftrightarrow\; \text{the } \vec{j}\text{-th } d\text{-cell of the } d \times n\text{-partition } \vec{x} \text{ contains an edge of } H$$

The premise of the theorem says that the sets $Y_{\vec{j}}$, $\vec{j} \in [n]^d$, cover $P_{d,n}$. It is easy to verify that the sets $Y_{\vec{j}}$ are open. Moreover, if any of the cells in a d-cell is empty then this d-cell cannot contain an edge of H. Therefore, for $\vec{x} \in P_{d,n}$, $i \in [d]$, and $\vec{j} \in [n]^d$ we have the implication $x_{j_i}^i = 0 \Rightarrow \vec{x} \notin Y_{\vec{j}}$.

We now associate with every $\vec{x} \in P_{d,n}$ a d-dimensional $n \times \cdots \times n$ array of nonnegative real numbers $S(\vec{x}) = (s_{\vec{j}}(\vec{x}))_{\vec{j} \in [n]^d}$, defined by:

$$s_{\vec{j}}(\vec{x}) = \operatorname{dist}(\vec{x}, Y_{\vec{j}}^c)$$

where $Y_{\vec{j}}^c = P_{d,n} \setminus Y_{\vec{j}}$ is a closed set, and "dist" denotes Euclidean distance. Note that for no $\vec{x} \in P_{d,n}$ do all the entries of $S(\vec{x})$ vanish, because that would mean that $\vec{x} \in Y_{\vec{j}}^c$ for all $\vec{j} \in [n]^d$, contradicting the fact that the sets $Y_{\vec{j}}$ cover $P_{d,n}$. Note also that $S(\vec{x})$ is continuous in $\vec{x}$.

Next, given the array $S(\vec{x})$ we denote by $\sigma_j^i(S(\vec{x}))$, for $i = 1, \ldots, d$, $j = 1, \ldots, n$, the sum of the entries in the j-th layer in direction i, that is:

$$\sigma_j^i(S(\vec{x})) = \sum_{\vec{j} \in [n]^d : j_i = j} s_{\vec{j}}(\vec{x})$$

Then we define $A_j^i \subseteq P_{d,n}$ for $i = 1, \ldots, d$, $j = 1, \ldots, n$, by:

$$\vec{x} \in A_j^i \Leftrightarrow \sigma_j^i(S(\vec{x})) = \max_{j' \in [n]} \sigma_{j'}^i(S(\vec{x}))$$

That is, $\vec{x} \in A_j^i$ if the j-th layer in direction i of $S(\vec{x})$ has the largest sum among the layers parallel to it.

We are going to apply Corollary 3.2 to the sets A_j^i. By the continuity of $S(\vec{x})$, these sets are closed. For any direction i, one of the layers has the largest sum, hence A_j^i, $j = 1, \ldots, n$, cover $P_{d,n}$. Finally, let us prove the implication $x_j^i = 0 \Rightarrow \vec{x} \notin A_j^i$. Assume, for contradiction, that $x_j^i = 0$ and $\vec{x} \in A_j^i$. As mentioned above, we deduce from $x_j^i = 0$ that for any $\vec{j} \in [n]^d$ with $j_i = j$ we have $\vec{x} \notin Y_{\vec{j}}$. Thus $\sigma_j^i(S(\vec{x})) = 0$, and we deduce from $\vec{x} \in A_j^i$ that $\sigma_{j'}^i(S(\vec{x})) = 0$ for all $j' \in [n]$, meaning that all the entries of $S(\vec{x})$ vanish. As shown above, this is impossible.

By Corollary 3.2, we find a point $\vec{x} \in \bigcap_{i=1}^d \bigcap_{j=1}^n A_j^i$. This means that in every direction, all layer-sums of $S(\vec{x})$ are equal. Hence, $\sigma_j^i(S(\vec{x})) = a$ for some fixed a, regardless of i and j. Clearly $a > 0$, as $S(\vec{x})$ cannot entirely vanish.

We now construct an auxiliary d-partite hypergraph Γ. Its vertex set $V(\Gamma)$ is partitioned into $V^i = \{v_1^i, \ldots, v_n^i\}$, $i = 1, \ldots, d$. We think of vertex v_j^i as representing the j-th cell in the i-th n-partition corresponding to the point $\vec{x}$ found in the previous paragraph. The edge set $E(\Gamma)$ consists of those sets $e_{\vec{j}} = \{v_{j_1}^1, \ldots, v_{j_d}^d\}$ such that $s_{\vec{j}}(\vec{x}) > 0$. That is, an edge indexed by $\vec{j}$ is present in Γ if $\vec{x} \in Y_{\vec{j}}$, i.e., the $\vec{j}$-th d-cell of $\vec{x}$ contains an edge of H.

The function $f : E(\Gamma) \to \mathbb{R}_+$ defined by $f(e_{\vec{j}}) = \frac{s_{\vec{j}}(\vec{x})}{a}$ is a fractional matching in Γ satisfying $\sum_{e \in E(\Gamma)} f(e) = n$. It follows from Theorem 3.4 that there is a matching M in Γ of size at least $\frac{n}{d-1}$. For every $e_{\vec{j}} \in M$, the corresponding $\vec{j}$-th d-cell of $\vec{x}$ contains an edge of H, yielding the conclusion of Theorem 1.5. □

Proof of Theorem 1.6: Given the d-interval hypergraph $H = (V, E)$ and the positive integer n, we define $Z_{\vec{j}} \subseteq P_{d,n}$ for $\vec{j} \in [n]^d$ by:

$\vec{x} \in Z_{\vec{j}} \Leftrightarrow$ the $\vec{j}$-th d-cell of the $d \times n$-partition $\vec{x}$ is contained in an edge of H

The premise of the theorem says that the sets $Z_{\vec{j}}$, $\vec{j} \in [n]^d$, cover $P_{d,n}$. It is easy to check that the sets $Z_{\vec{j}}$ are closed.

For a fixed $\varepsilon > 0$, we associate with every $\vec{x} \in P_{d,n}$ a d-dimensional $n \times \cdots \times n$ array of nonnegative real numbers $T(\vec{x}, \varepsilon) = (t_{\vec{j}}(\vec{x}, \varepsilon))_{\vec{j} \in [n]^d}$, defined by:

$$t_{\vec{j}}(\vec{x}, \varepsilon) = \max(1 - \frac{\operatorname{dist}(\vec{x}, Z_{\vec{j}})}{\varepsilon}, 0)$$

For no $\vec{x} \in P_{d,n}$ do all the entries of $T(\vec{x}, \varepsilon)$ vanish, because that would mean that $\vec{x} \notin Z_{\vec{j}}$ for all $\vec{j} \in [n]^d$, contradicting the fact that the sets $Z_{\vec{j}}$ cover $P_{d,n}$. Note also that $T(\vec{x}, \varepsilon)$ is continuous in $\vec{x}$.

As in the previous proof, we consider for $i = 1, \ldots, d$, $j = 1, \ldots, n$, the sum of the entries in the j-th layer in direction i of the array $T(\vec{x}, \varepsilon)$, that is:

$$\sigma^i_j(T(\vec{x}, \varepsilon)) = \sum_{\vec{j} \in [n]^d : j_i = j} t_{\vec{j}}(\vec{x}, \varepsilon)$$

Then we define $B^i_j(\varepsilon) \subseteq P_{d,n}$ for $i = 1, \ldots, d$, $j = 1, \ldots, n$, by:

$$\vec{x} \in B^i_j(\varepsilon) \Leftrightarrow \sigma^i_j(T(\vec{x}, \varepsilon)) = \max_{j' \in [n]} \sigma^i_{j'}(T(\vec{x}, \varepsilon))$$

We are going to apply Theorem 3.3 to the sets $B^i_j(\varepsilon)$. The closedness and covering properties are verified just as in the previous proof. To check condition (b) of the theorem, we need the following:

Claim 4.1. *Let $\vec{x} \in P_{d,n}$ and $i \in [d]$. Consider $\vec{j}, \vec{k} \in [n]^d$ such that $j_{i'} = k_{i'}$ for all $i' \in [d] \setminus \{i\}$, and $x^i_{j_i} = 0$. Then $t_{\vec{j}}(\vec{x}, \varepsilon) \geq t_{\vec{k}}(\vec{x}, \varepsilon)$.*

Proof. By the definition of $T(\vec{x}, \varepsilon)$, it suffices to show that

$$\operatorname{dist}(\vec{x}, Z_{\vec{j}}) \leq \operatorname{dist}(\vec{x}, Z_{\vec{k}}).$$

Let $\vec{y} \in Z_{\vec{k}}$ be closest to $\vec{x}$, and let e be an edge of H that contains the $\vec{k}$-th d-cell of the $d \times n$-partition $\vec{y}$ (by the definition of $Z_{\vec{k}}$ such an edge exists). Consider $\vec{z} \in P_{d,n}$ which coincides with $\vec{x}$ in the i-th copy of Δ_{n-1}, and with $\vec{y}$ in all other copies. Then $\operatorname{dist}(\vec{x}, \vec{z}) \leq \operatorname{dist}(\vec{x}, \vec{y})$. Moreover, the $\vec{j}$-th d-cell of $\vec{z}$ consists of an empty cell in U^i (because $z^i_{j_i} = x^i_{j_i} = 0$), and otherwise coincides with the $\vec{k}$-th d-cell of $\vec{y}$, implying that the $\vec{j}$-th d-cell of $\vec{z}$ is also contained in e. This shows that $\vec{z} \in Z_{\vec{j}}$, and hence

$$\operatorname{dist}(\vec{x}, Z_{\vec{j}}) \leq \operatorname{dist}(\vec{x}, \vec{z}) \leq \operatorname{dist}(\vec{x}, \vec{y}) = \operatorname{dist}(\vec{x}, Z_{\vec{k}}).$$

□

Returning to the main proof, we verify the implication $x^i_j = 0 \Rightarrow \vec{x} \in B^i_j(\varepsilon)$. Indeed, assuming $x^i_j = 0$, when we compare $\sigma^i_j(T(\vec{x}, \varepsilon))$ and $\sigma^i_{j'}(T(\vec{x}, \varepsilon))$, we find by Claim 4.1 that each term in the former is at least as large as the corresponding term in the latter. This implies that $\sigma^i_j(T(\vec{x}, \varepsilon)) \geq \sigma^i_{j'}(T(\vec{x}, \varepsilon))$ for all $j' \in [n]$, thus $\vec{x} \in B^i_j(\varepsilon)$.

By Theorem 3.3, we find a point $\vec{x}(\varepsilon) \in \bigcap_{i=1}^{d} \bigcap_{j=1}^{n} B_j^i(\varepsilon)$. As before, this means that $\sigma_j^i(T(\vec{x}(\varepsilon), \varepsilon)) = a(\varepsilon)$ for some fixed $a(\varepsilon) > 0$, regardless of i and j.

We construct an auxiliary d-partite hypergraph $\Gamma(\varepsilon)$ in a way similar to the previous proof. The vertex set $V(\Gamma(\varepsilon)) = V(\Gamma)$ is partitioned into $V^i = \{v_1^i, \ldots, v_n^i\}$, $i = 1, \ldots, d$. The edge set $E(\Gamma(\varepsilon))$ consists of those sets $e_{\vec{j}} = \{v_{j_1}^1, \ldots, v_{j_d}^d\}$ such that $t_{\vec{j}}(\vec{x}(\varepsilon), \varepsilon) > 0$. That is, an edge indexed by $\vec{j}$ is present in $\Gamma(\varepsilon)$ if $\operatorname{dist}(\vec{x}(\varepsilon), Z_{\vec{j}}) < \varepsilon$.

The function $g : E(\Gamma(\varepsilon)) \to \mathbb{R}_+$ defined by $g(e_{\vec{j}}) = \frac{t_{\vec{j}}(\vec{x}(\varepsilon), \varepsilon)}{a(\varepsilon)}$ is a fractional edge-cover in $\Gamma(\varepsilon)$ satisfying $\sum_{e \in E(\Gamma(\varepsilon))} g(e) = n$. It follows from Corollary 3.6 that there is an edge-cover $C(\varepsilon)$ in $\Gamma(\varepsilon)$ of size at most $(1 + \ln d)n$.

Up till now, $\varepsilon > 0$ was fixed. Because $P_{d,n}$ is compact and the number of possibilities for the edge-cover $C(\varepsilon)$ is finite, we can find a sequence of positive numbers $\varepsilon_k \to 0$ such that $\vec{x}(\varepsilon_k)$ converges to some $\vec{x} \in P_{d,n}$, and $C(\varepsilon_k)$ is a constant set C of at most $(1 + \ln d)n$ d-tuples that together cover $V(\Gamma)$. For any $e_{\vec{j}} \in C$ we have $\operatorname{dist}(\vec{x}(\varepsilon_k), Z_{\vec{j}}) < \varepsilon_k$ for all k, implying that $\vec{x} \in Z_{\vec{j}}$. Thus, for any $e_{\vec{j}} \in C$, the corresponding $\vec{j}$-th d-cell of $\vec{x}$ is contained in an edge of H. As the $e_{\vec{j}} \in C$ cover $V(\Gamma)$, we have at most $(1 + \ln d)n$ edges of H that together cover all cells of $\vec{x}$, and hence cover V, as required.

In the special case $d = 2$, the hypergraphs $\Gamma(\varepsilon)$ are actually bipartite graphs, and so we may use Theorem 3.7 instead of Corollary 3.6, and get n as an upper bound instead of $(1 + \ln d)n$. □

5. Are the bounds tight?

The $\frac{n}{d-1}$ bound in Theorem 1.5 is obviously tight for $d = 2$. For large d, it has been shown by Matoušek [11] that the bound is tight up to a factor of order $\log^2 d$.

Regarding Theorem 1.6, the following simple example shows that the bound $\rho(H) \leq n$ for $d = 2$ is best possible.

Example 5.1. Let $V = U^1 \cup U^2$, and let n be a positive integer. Consider the 2-interval hypergraph $H = (V, E)$ having n^2 edges of the form $e = I_k^1 \cup I_\ell^2$, $k = 1, \ldots, n$, $\ell = 1, \ldots, n$, where $I_j^i = [\frac{j-1}{n}, \frac{j}{n}] \subseteq U^i$. By the pigeonhole principle, every $2 \times n$-partition of V has a 2-cell that is contained in an edge of H. Clearly, $\rho(H) = n$. ◊

Moving to $d = 3$, the following example shows that $\rho(H) \leq n$ cannot be guaranteed anymore.

Example 5.2. Let $V = U^1 \cup U^2 \cup U^3$, and let $n = 2$. For $i = 1, 2, 3$, $v = 0, 1$, and $0 \leq \ell < 1$, we denote by $I_{v,\ell}^i$ the closed subinterval of U^i that contains its endpoint v and has length ℓ. We construct a 3-interval hypergraph $H = (V, E)$ in which every edge is indexed by some $\vec{v} \in \{0, 1\}^3$ and some

$\vec{\ell} \in [0,1)^3$, and is of the form:

$$e_{\vec{v},\vec{\ell}} = I^1_{v_1,\ell_1} \cup I^2_{v_2,\ell_2} \cup I^3_{v_3,\ell_3}$$

We put $e_{\vec{v},\vec{\ell}}$ in E only for certain choices of $\vec{v}, \vec{\ell}$, namely those satisfying:

- $\ell_i \in \{\frac{k}{24} \mid k = 0, 1, \ldots, 23\}$, $i = 1, 2, 3$

and either

- $\vec{v} \in \{(0,0,0), (1,0,0), (0,1,0), (0,0,1)\}$ and $\ell_1 + \ell_2 + \ell_3 = \frac{47}{24}$

or

- $\vec{v} \in \{(1,1,1), (0,1,1), (1,0,1), (1,1,0)\}$ and $\ell_1 + \ell_2 + \ell_3 = 1$

We proceed to show that the premise of Theorem 1.6 is satisfied for $n = 2$. Namely, that every 3×2-partition of V has a 3-cell that is contained in an edge of H. Let

$$[0, c^1), (c^1, 1]; \; [0, c^2), (c^2, 1]; \; [0, c^3), (c^3, 1]$$

be such a 3×2-partition. First, assume that $\frac{1}{24} \leq c^i \leq \frac{23}{24}$ for $i = 1, 2, 3$. If $c^1 + c^2 + c^3 \leq \frac{15}{8}$, then rounding each c^i up to the nearest number of the form $\frac{k}{24}$, the sum will not exceed $\frac{47}{24}$, and hence the 3-cell $[0, c^1) \cup [0, c^2) \cup [0, c^3)$ is contained in some $e_{\vec{v},\vec{\ell}} \in E$ with $\vec{v} = (0, 0, 0)$. Similarly, if $(1 - c^1) + (1 - c^2) + (1 - c^3) \leq \frac{7}{8}$, the 3-cell $(c^1, 1] \cup (c^2, 1] \cup (c^3, 1]$ is contained in some $e_{\vec{v},\vec{\ell}} \in E$ with $\vec{v} = (1, 1, 1)$. In either case we are done, therefore we may assume that:

$$\frac{15}{8} < c^1 + c^2 + c^3 < \frac{17}{8}$$

Arguing similarly with respect to each of the other three pairs of antipodal $\vec{v} \in \{0,1\}^3$, we see that we are done unless the following three constraints hold as well:

$$\frac{15}{8} < (1 - c^1) + c^2 + c^3 < \frac{17}{8}$$
$$\frac{15}{8} < c^1 + (1 - c^2) + c^3 < \frac{17}{8}$$
$$\frac{15}{8} < c^1 + c^2 + (1 - c^3) < \frac{17}{8}$$

Combining each of these three constraints with the first one above yields:

$$|c^i - (1 - c^i)| < \frac{1}{4}, \; i = 1, 2, 3$$

This, in turn, yields $\frac{3}{8} < c^i < \frac{5}{8}$, $i = 1, 2, 3$, contradicting $c^1 + c^2 + c^3 > \frac{15}{8}$.

It remains to handle cases where one or more cells in the given 3×2-partition are shorter than $\frac{1}{24}$ ("short" for brevity). Obviously, there is at most one short cell in each U^i. If there are short cells in all U^i, then the 3-cell that they form is easily contained in an edge of H. If there are short cells in two of the U^i, we look at the 3-cell that they form together with the shorter of the two cells in the third U^i; again, it is easily contained in an edge of H. Finally, suppose that there is a short cell in only one of the

U^i, say U^1. If the short cell is $[0, c^1)$, then the 3-cell $[0, c^1) \cup [0, c^2) \cup [0, c^3)$ is contained in $e_{(0,0,0),(\frac{1}{24},\frac{23}{24},\frac{23}{24})}$. If the short cell is $(c^1, 1]$ then the 3-cell $(c^1, 1] \cup [0, c^2) \cup [0, c^3)$ is contained in $e_{(1,0,0),(\frac{1}{24},\frac{23}{24},\frac{23}{24})}$. This completes the verification of the premise of Theorem 1.6 for $n = 2$.

On the other hand, $\rho(H) = 3$ in this example. To see that there is no edge-cover of size 2, assume to the contrary that $e_{\vec{v},\vec{\ell}}, e_{\vec{v'},\vec{\ell'}} \in E$ cover V. Then necessarily $\vec{v}, \vec{v'}$ are antipodal in $\{0,1\}^3$, and therefore, of the two sums $\ell_1 + \ell_2 + \ell_3$ and $\ell'_1 + \ell'_2 + \ell'_3$, one equals $\frac{47}{24}$ and the other 1. Hence the total length of these two 3-intervals is $\frac{71}{24}$, insufficient to cover V. $\Diamond$

The above example suggests that for $d = 3$ it might be possible to improve the bound $\rho(H) \leq (1 + \ln 3)n$ to $\rho(H) \leq \frac{3}{2}n$. While we are currently unable to achieve this, we can reduce the constant somewhat, and prove $\rho(H) < \frac{7}{4}n$. This follows from the respective improvement of the bound of Corollary 3.6 under the conditions relevant to our application.

To state this improvement, we recall that a function $f : E \to \mathbb{R}_+$ is called a *perfect fractional matching* (or equivalently a *perfect fractional edge-cover*) in $H = (V, E)$ if $\sum_{e \ni v} f(e) = 1$ for all $v \in V$.

Proposition 5.3. *Let $H = (V, E)$ be a 3-partite hypergraph, with $V = V^1 \cup V^2 \cup V^3$ where $|V^i| = n$ for $i = 1, 2, 3$. If H has a perfect fractional matching, then $\rho(H) < \frac{7}{4}n$.*

Proof. Let $f : E \to \mathbb{R}_+$ be a perfect fractional matching in H. Using only edges in the support of f, we construct an edge-cover C in three steps:

Step 1: We find a matching M of maximal size, and place its edges in C.

Step 2: For as long as we can find an edge covering two of the yet uncovered vertices, we add such edges to C.

Step 3: We cover the remaining uncovered vertices one-by-one, adding suitable edges to C.

By Theorem 3.4, we have $|M| \geq \frac{n}{2}$ for the matching in Step 1. We also observe that the set V_0 of uncovered vertices after the completion of Step 2 satisfies $|V_0| < n$. Indeed, the edges that intersect V_0 do so in only one vertex, and therefore contribute twice as much f-weight to $V \setminus V_0$ than to V_0. Since the total f-weight on every vertex is 1, this implies that $2|V_0| \leq |V \setminus V_0|$, or $|V_0| \leq n$. The strict inequality comes from the fact that there is positive contribution of f-weight to $V \setminus V_0$ from edges used in the first two steps (recall that they are in the support of f).

Now we can estimate the number of edges used as follows:

$$|C| = |M| + \frac{1}{2}(3n - 3|M| - |V_0|) + |V_0| = \frac{1}{2}(3n - |M| + |V_0|) < \frac{1}{2}(3n - \frac{n}{2} + n) = \frac{7}{4}n$$

□

Corollary 5.4. *Under the condition of Theorem 1.6 for $d = 3$, we have $\rho(H) < \frac{7}{4}n$.*

An analogous improvement of the bound in Theorem 1.6, taking advantage of d-partiteness, can be obtained for arbitrary d. But it only leads to an insignificant improvement of the constant $1 + \ln d$ for large d, hence we do not state it explicitly. Unfortunately, we do not know if the bound $\rho(H) \leq (1 + \ln d)n$ is close to best possible as d grows to infinity.

References

[1] R. Aharoni, T. Kaiser, and S. Zerbib, Fractional covers and matchings in families of weighted d-intervals, to appear in Combinatorica, arXiv:1402.2064v2.

[2] N. Alon, Piercing d-intervals, Disc. Comput. Geom. 19 (1998), 333–334.

[3] E. Berger, KKM — A topological approach for trees, Combinatorica 25 (2004), 1–18.

[4] Z. Füredi, Maximum degree and fractional matchings in uniform hypergraphs, Combinatorica 1 (1981), 155–162.

[5] T. Gallai, Über extreme Punkt-und Kantenmengen (in German), Ann. Univ. Sci. Budapest, Eötvös Sect. Math. 2 (1959), 133–138.

[6] T. Gallai, Graphen mit triangulierbaren ungeraden Vielecken (in German), Magyar Tud. Ak. Mat. Kut. Int. Közl. 7 (1962), 3–36.

[7] A. Gyárfás and J. Lehel, Covering and coloring problems for relatives of intervals, Discrete Mathematics 55 (1985), 167–180.

[8] T. Kaiser, Transversals of d-intervals, Disc. Comput. Geom. 18 (1997), 195–203.

[9] B. Knaster, C. Kuratowski, and S. Mazurkiewicz, Ein Beweis des Fixpunktsatzes für n-dimensionale Simplexe (in German), Fundamenta Mathematicae 14 (1929), 132–137.

[10] L. Lovaśz, On the ratio of optimal integral and fractional covers, Discrete Math. 13 (1975), 383–390.

[11] J. Matoušek, Lower bounds on the transversal numbers of d-intervals, Disc. Comput. Geom. 26 (2001), 283–287.

[12] B. Peleg, Equilibrium points for open acyclic relations, Canad. J. Math. 19 (1967), 366–369.

[13] B. Peleg, Existence theorem for the bargaining set $M_1^{(i)}$, in: M. Shubik, ed., Essays in mathematical economics in honor of Oskar Morgenstern (Princeton University Press, Princeton, 1967), 53–56.

[14] E. Sperner, Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes (in German), Abh. Math. Sem. Univ. Hamburg 6 (1928), 265–272.

[15] G. Tardos, Transversals of 2-intervals, a topological approach, Combinatorica 15 (1995), 123–134.

Department of Mathematics, Technion, Haifa, Israel
E-mail address: raharoni@gmail.com

Department of Mathematics, Technion, Haifa, Israel
E-mail address: holzman@tx.technion.ac.il

Department of Mathematics, University of Michigan, Ann Arbor, USA
E-mail address: zerbib@umich.edu